



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/505,807	02/17/2000	Abraham P. Ittycheriah	YO999-195 (8728-281)	7313
7590	07/08/2004		EXAMINER	
Frank Chau Esq F Chau & Associates LLP 1900 Hempstead Turnpike Suite 501 East Meadow, NY 11554			FOSTER, ROLAND G	
			ART UNIT	PAPER NUMBER
			2645	
			DATE MAILED: 07/08/2004	

Please find below and/or attached an Office communication concerning this application or proceeding.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

MAILED

JUL 08 2004

Technology Center 2600

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Paper No. 12

Application Number: 09/505,807
Filing Date: February 17, 2000
Appellant(s): ITTYCERIAH ET AL.

Frank DeRosa
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed March 31, 2004.

(1) *Real Party in Interest*

A statement identifying the real party in interest is contained in the brief.

(2) *Related Appeals and Interferences*

A statement identifying the related appeals and interferences which will directly affect or be directly affected by or have a bearing on the decision in the pending appeal is contained in the brief.

(3) *Status of Claims*

The statement of the status of the claims contained in the brief is correct.

(4) *Status of Amendments After Final*

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) *Summary of Invention*

The summary of invention contained in the brief is correct.

(6) *Issues*

The appellant's statement of the issues in the brief is correct.

(7) *Grouping of Claims*

Appellant's brief includes a statement that dependent claims 3-6, 8, 17, 18, 20, 24, 25, and 27 do not stand or fall together and provides reasons as set forth in 37 CFR 1.192(c)(7) and (c)(8). For example, applicant separately argues the patentability of the subject dependent claims in the argument section (8).

(8) *ClaimsAppealed*

The copy of the appealed claims contained in the Appendix to the brief is correct

(9) *Prior Art of Record*

6,519,686 B2	Woodring et al.	2-2003
4,916,658	Lee et al.	4-1990

(10) *Grounds of Rejection*

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 102

Claims 1, 3-6, 8-18, 20-25, and 27 are rejected under 35 U.S.C. 102(e) as being anticipated by U.S. Patent No. 6,519,686 B2 to Woodring et al. (hereinafter "Woodring"), of record.

With respect to claim 1, see the following paragraphs for details on how Woodring anticipates particular limitations within the claim.

The limitation "a first queue for storing data" reads on the abstract, Fig. 4, and col. 7, lines 19-67, where the buffer 372 queues data originating from producer 310 for transmission to consumers 0 through N-1.

The limitation "a data source for outputting the data that is stored in the first queue" reads on Fig. 4, producer 310 that outputs data for storage into a buffer in buffer storage 372 (first queue).

The limitation "a plurality of consumers each sharing the data stored in the first queue" reads on the abstract, Fig. 4, and col. 6, lines 1-32, where a plurality of consumers 0 through N-1 share data stored in buffer 372 (first queue) that was generated by producer 310 as discussed above.

The limitation "a scheduler for managing the storage and consumption of the data in the first queue and for controlling the data source and the plurality of consumers to control the amount of data stored in and consumed from the first queue" reads on Woodring as follows. The storage manager 350 manages the storage of data as its name implies including the data stored in the buffer 372 (first queue). The storage manager 350 also stores information related to the each consumers desired sample notification rate. The sample notification rate relates to the amount of

data being streamed to each client (col. 8, lines 65-67). The amount of data streamed to each client is consumption data in the first queue by each of the plurality of consumers. Therefore, the storage manager 350 reads on the scheduler. Woodring discloses that the storage manager 350 includes a free buffer semaphore (FBSEM) mechanism 376 (Fig. 4). The FBSEM 376 controls the producer 310, for example, by signaling the producer that one or more buffers in buffer storage 372 (first queue) are free (col. 7, lines 61-63). The storage manager 350 also includes the producer mutual exclusion mechanism (MUTEX) 374, which controls the producer by placing a requirement on the producer to acquire exclusive ownership (col. 7, lines 52-60) before accessing management data for the buffer storage (first queue). The storage manager 350 also includes a variety of mechanisms to control the clients (consumers) as they receive data from the buffer storage 372 (first queue) such as the mail slots (col. 7, lines 1-20), management data structure and information streams (col. 7, lines 19-39), and the buffer masks (col. 7, lines 45-52).

Claim 16 differs substantively from claim 1 in that claim 16 recites a method whose steps performs functions equivalent to the system components recited in claim 1. Therefore, see the claim 1 rejection for further details.

Claim 23 differs substantively from claim 16 in that claim 23 performs a method equivalent to the method of claim 16. Therefore, see the claim 16 rejection for additional details. In addition, the claim 23 method is performed using program instructions stored on a program storage device and executed by a machine. The queuing method of Woodring is implemented on

a computer (Fig. 1) which implements the method using program instruction stored on a program storage device (e.g., memory) and excused by a machine (e.g., a processor). For example, see col. 3, line 61 – col. 4, line 5.

With respect to claim 3, the producer 310 (data source) comprises PCs 11 and 12 engaged in a video conference (Fig. 1).¹ During the conference, the PC 12 is capable of both transmitting and receiving data (col. 2, lines 57-60). Therefore, PC 12 (comprising the data source) is also a consumer of multimedia streaming data that it receives from the queue of PC 11 (a second queue).²

With respect to claims 4-6, 17, 18, 24, and 25, see col. 8, line 50 – col. 9, line 5. The priority request relates to the desired sample rate and to the key attribute.³

With respect to claims 8, 20, and 27, each buffer holds an amount of data represented by the buffer size 368 (col. 7, lines 32-35). All consumers allocated to those buffers read out the

¹ Note that producer 310 is part of Fig. 4., where Fig. 4 "is generic to any producer-consumer relationship, whether it is local or remote (col. 6, lines 8-17) (emphasis added). Thus, producer 310 would be generic to the producer-consumer relationship illustrated in the video-conferencing arrangement of Fig. 1, where PC 12 (part of producer 310) would be associated with a first queue for locally produced data and a second queue for consuming remotely produced data from PC 11. Otherwise if producer 310 was insufficient to handle the duties illustrated in Fig. 1, then Fig. 4 would not be generic to any producer-consumer relationship whether local or remote as required by Woodring.

² Similarly to producer 310 (see footnote 1), the storage manager 350 (scheduler) of Fig. 4 is generic to Fig. 1. Also as previously discussed (see the claim 1 rejection), the storage manager 350 (scheduler) controls the producer-consumer queues. Therefore, the storage manager 350 would of course control the first and second queues discussed above.

³ Note that examiner's interpretation of priority is consistent with applicant's specification. For example, applicant specification states that one example of "prioritizing and slowing down different consumers" is to "slow down network traffic towards fast consumers" (page 27, lines 10-20). Slowing down network traffic to the consumers relates directly to Woodring's sampling rate, which controls the speed (rate) in which the consumers receive data from the network.

amount of data stored in the buffer (col. 7, lines 13-18). Only when the last consumer releases from the buffer after reading will the buffer mask 364 be set to zero thereby allowing the buffer to become available for other data consumption uses (e.g., different consumers). Therefore, the storage manager 350 (scheduler) prioritizes data consumption by keeping buffers allocated for at least one consumer if the amount of data stored in those buffers is not completely read out (i.e., unread) by that one consumer.

With respect to claims 9 and 21, the system's "engine" (e.g., the processors) are embedded within user stations 11 and 12 as discussed above.

With respect to claims 10 and 22, see col. 2, line 62.

With respect to claim 11, see col. 5, line 11.

With respect to claim 12, the data traffic management that the storage manager 350 (scheduler) performs for each consumer as discussed above represents data traffic that is sent out over the network. Therefore, the scheduler monitors, manages, and schedules data that is sent over the network.

With respect to claim 13, the system is dynamically re-programmed each time a new consumer attaches and registers (col. 8, lines 50-67) as discussed above.

With respect to claim 14, attachment of consumers represents the attachment of video conferencing engine that supports both transmit and receive functions so that all the consumers can engage in a video conversation. The attachment process itself also includes conversational engines and functions as discussed above. Finally, the mail slot system also allows the consumers to engage conversationally with the system (col. 7, lines 1-20).

With respect to claim 15, see col. 8, lines 19-22.

Claim Rejections - 35 USC § 103

Claims 7, 19, and 26 are rejected under 35 U.S.C. 103(a) as being unpatentable over Woodring as applied to claims 1, 2, 16, 17, 23, and 24 above, and further in view of U.S. Patent No. 4,916,658 to Lee et al. (hereinafter "Lee"), of record.

Although Woodring discloses that the scheduler maintains a queue comprising buffers for each consumer as discussed above, Wood fails to disclose maintaining an IN pointer for the data source and an OUT pointer for each consumer.

However, Lee teaches that a commonly used type of buffer is a circular buffer which contains an IN pointer for the source of data and an OUT pointer for the consumer of data (col. 1, lines 10-41).

Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to add a circular buffer (comprising an IN pointer for the source of data and an OUT pointer) to the buffer for each consumer as disclosed by Woodring. This obvious addition of an OUT pointer to each buffer results in OUT pointer for each consumer because there is a buffer for each consumer as discussed above.

The suggestion/motivation for doing so would have been that circular buffers are "quite commonly used in data processing systems" (Lee, col. 1, lines 24-26). In addition, buffer storage flexibility would have been increased because a "circular queue may contain any number of storage positions" (Lee, col. 1, lines 28-31). However, the "management of the circular queue requires that the 'head' [IN pointer] and 'tail' [OUT pointer] of the queue must be known" (Lee, col. 29-34). Therefore, the obvious addition of just an operational circular queue would have required the use of both an IN and OUT pointer.

(11) Response to Argument

Preliminary Comparison of Applicant's Invention to Woodring

Both applicant's invention and Woodring are directed toward streaming data (e.g., audio streams) from a data source to plural consumers. For example, applicant's invention splits an audio stream from an audio subsystem 300 (data source) stored in a queue 301 into multiple consumer streams for consumption by multiple consumers 1-N using a scheduler 302 (e.g., Fig. 5). Similarly, Woodring splits an audio (and/or video) stream from producer 310 (data source)

stored in buffer storage 372 (queue) into multiple "consume" streams for consumption by multiple consumers 3200 to 320N-1 using a storage manager 350 (scheduler) (Fig. 4).⁴

Basic Issue

Applicant's basic argument (e.g., see pages 6 and 7 of the brief) is that the scheduler (storage manager 350) in Woodring fails actively "control" the data source (producer 310) and the consumers to, in turn, "control" the amount of data stored in and consumed from the queue (a particular buffer in buffer storage 372) (as required by the claims).

Applicant's argument is not persuasive.

First, the limitation "actively control" does not appear in the claim. It is not persuasive for the applicant to argue limitations when those limitations do not appear in the claim.

Second and nonetheless, Woodring's scheduler (storage manager 350) comprises an extensive management structure (Fig. 4) that places active requirements on (i.e., actively controls) how the data source and consumer go about controlling the amount of data stored in and consumed from the scheduler's own queue (buffer storage 372) as discussed in the last Office action and again below. If it were otherwise, the storage manager (scheduler) would not be managing how other devices store data onto its buffer, despite the name storage manager.

⁴ Note that the broadcast mailslots 0 – N-1 are NOT audio/video streams but are instead control channels such as those mailslots used in applicant's invention (see Fig. 1, mailbox 14).

Third, the recited term "control" should be interpreted reasonably broad, which is not consistent with equating the term to some undefined, unclaimed, and arbitrary level of active or supervisory control sufficient to distinguish over the prior art. By such a narrow interpretation, entirely passive and well known controllers such as Proportional-Integral-Derivative ("PID") controllers, comprising passive components such as capacitors and inductors, would not be considered as performing control functions, which they clearly are.

An examiner has the duty to give claim language its broadest reasonable interpretation so that the applicant fashions precise and clear claims by amending the claims, where amending the claims under such reasonable circumstances does not result in any unfairness or impairment to the applicant.⁵ Here, the term "control" is a very broad, commonplace, and non-technical term with no special meaning and whose ordinary and plain meaning would certainly include exercising a restraining or directing influence as consistent with the examiner's interpretation in the last Office action.⁶ Thus, it would have been helpful to the public and fair to all if the applicant had introduced express language more precisely and clearly defining what type of

⁵ "[A]n examiner has the duty to police claim language by giving it the broadest reasonable interpretation." Springs Window Fashions LP v. Novo Industries, L.P., 65 USPQ2d 1826, 1830 (Fed. Cir. 2003). "An essential purpose of patent examination is fashion claims that are precise, clear, correct and unambiguous. Only in this way can uncertainties of claim scope be removed, as much as possible, during the administrative process." In re Zletz, 13 USPQ2d 1320, 1322 (Fed. Cir. 1989). "Construing claims broadly during prosecution is not unfair to the applicant...because the applicant has the opportunity to amend the claims to obtain more precise claim coverage." In re American Academy of Science Tech Center, 03-1531, *7 (Fed. Cir., May 13, 2004) (Fed. Cir. BBS). "Applicants' interests are not impaired since they are not foreclosed from obtaining appropriate coverage for their invention with express claim language." In re Yamamoto, 740 F.2d 1569, 1571-72 (Fed. Cir. 1984).

⁶ Examiner's plain meaning interpretation of "control" would have also been well known to one of ordinary skill, as evidenced by the dictionary definition of the term. "The first step in claim construction is to determine the ordinary and customary meaning...Dictionary definitions frequently are useful in this process...." Boehringer Ingelheim Vetmedica, Inc. v. Schering-Plough Corp., 65 USPQ2d 1961, 1965 (Fed. Cir. 2003). For example, the verb

"control" is taking place, especially since applicant has chosen to even more broadly multiply the term "control" upon itself (e.g., A "controls" B "controls" C) (e.g., see claim 1). A fair reading of Woodring leaves no doubt that scheduler (storage manager 350) exercises a restraining or directing influence on (i.e., controls) both the data source (producer 310) and each of the consumers 320₀ to 320_{N-1}, and that they, in turn, exercise a restraining or directing influence on (i.e., control) the amount of data stored in and consumed from the queue (a particular buffer within buffer storage 372), as discussed in the prior Office action and below.

Specific Applicant Arguments

Claims 1, 16, and 23

Regarding specific applicant arguments, on page 6 of the brief, the applicant argues that Woodring "expressly teaches that synchronization, management and processing of the information stream between the producer (310) and the consumers (320) is achieved via the producer-side interface (410) and a consumer-side interface (450) (see, Fig. 5), which comprise software mechanisms that are implemented as handles to objects that are shared by the producer and consumers (see, e.g., Col. 6, lines 18-32; and Col. 8, lines 31 et seq.)."

Applicant's arguments are not persuasive. Applicant is overlooking the intermediate inter-process communication ("IPC") channel 330 that stands between the producer 310 (data source) and consumer interfaces 320. That is, the producer and consumer interfaces referred to by the applicant interface not with each other but directly with IPC 330, which comprises the

"control" is defined as "to exercise restraining or directing influence over." Merriam Webster's Collegiate

storage manager 350 (scheduler) (Fig. 4 and col. 6, lines 62-67). Thus, synchronization and management of the stream between the producer and consumers is achieved on the back of storage manager 350 (scheduler). Further, the existence of shared objects is not particularly relevant to the question of whether the scheduler controls the producer and consumer, and to the extent that it is, these shared objects (e.g., Mutex and FreeBuffSem) are actually management structures that reside in the storage manager 350 (Fig. 4). Thus, the management and synchronization objects for the producer 310 (data source) and consumer 320 are stored in the scheduler (storage manager 350).

On page 6 of the brief, the applicant argues that the "storage manager (350) simply stores the management data structure that is used by the producer and consumer for purposes of communication." Thus, the storage manager 350 (scheduler) fails to actively control the data source and consumer to control the amount of data in the queue (buffer).

Applicant's arguments are not persuasive. Applicant admits above that the scheduler (storage manager 350) stores the management data structure used by the producer 310 (data source) and consumer 320 for the purposes of communication. Thus, applicant apparently argues that the scheduler comprises management data structure for both the producer (data source) and consumers yet somehow fails to exercise a restraining or directing influence on (i.e., controls) them in some manner. It would seem improbable that a device would contain

management data structure yet actually not control in some manner. Nonetheless, examples of control are discussed below in response to further applicant arguments.

On page 7 of the brief, the applicant argues that the "[e]xaminer has not explained how the FBSEM mechanism (376) actually 'controls the data source .. to control the amount of data stored in the queue' as essentially claimed." Instead, the applicant argues that the FBSEM 376 is "merely a notification mechanism that enables the producer to actually determine (via the producer BSEM mechanism (316), Fig. 4) what buffers (multiple queues) are...not available."

Although the applicant's argument that the scheduler must "essentially" control the amount of data stored in the queue via control of the data source is addressed in the next paragraph, it is questionable whether the claims require this. Specifically, the applicant expressly claims a "scheduler...for controlling the data source and the plurality of consumers to control the amount of data stored in and consumed from the first queue" (emphasis added). That is, the scheduler controls the data source and consumers, where the source and consumers in turn control the amount of data stored in and consumed from the queue. Thus, the scheduler controls the amount of data only in an indirect, and thus attenuated manner, if at all. For example, the data source may make a substantially independent decision to control the amount of data written to the queue after receiving control signals substantially directed to other subject matter from the scheduler. Such a example would squarely fall under the scheduler controlling the data source, which in turn (for substantially independent reasons) controls the amount of data written to the queue. Otherwise the second use of the term "control" would have to be overlooked entirely, or

at least de-emphasized. This is not a preferred method of claim interpretation, since the applicant could have simply and easily recited that the scheduler controlled the amount of data. It should be presumed that the applicant expressly recited and inserted the second use of the term "control" for a reason. Thus, there is insufficient basis in the claims to argue that the scheduler must "essentially" control the amount of data in the queue via control of the data source and consumer because the scheduler may only control the amount of data in an indirect and thus attenuated manner via the data source and consumer.

Nonetheless, applicant's arguments that the scheduler's (data source's) free buffer semaphore ("FBSEM") 376 fails to control the data source (producer 310), which in turn controls the amount of data in the queue, are unpersuasive. Indeed, the applicant seems to admit that the scheduler's FBSEM 376 controls the data source, which in turn controls the amount of data in the queue (buffer). Specifically, a notification mechanism that provides information regarding the availability of a buffer (queue) to the producer 310 (scheduler) would seem to be exercising a restraining or directing influence on (i.e., controlling) the producer (data source) to in turn exercise a restraining or directing influence on (i.e., control) the amount of data written by the producer to the queue (buffer). *For example, if the scheduler's FBSEM does not provide a notification to the producer (data source) that the buffer (queue) is available, then no amount of data will be written by the producer (data source) into the buffer (queue).* Nonetheless beyond applicant's admission above, the examiner also provided extensive support for the conclusion that the FBSEM 376 controlled the amount of data stored in a queue (buffer) via control of the producer 310 (data source), such as noting that the FBSEM signals the producer that one or more

buffers in the buffer storage (queue) are free for data storage (see the rejection above, page 6 of the last Office action, or col. 7, lines 61-63).

On page 7 of the brief, the applicant argues that the MUTEX 374 does not control the producer because the MUTEX does nothing more than prevent conflicting accesses by providing mutually exclusive access by the producer (data source) and consumers to the management data structure with the exception of the buffer storage area. Thus, the MUTEX 374 cannot control the producer (data source) to control the amount of data stored in a queue. Further, even though the examiner explained that the MUTEX 374 controls the producer by placing a requirement on the producer to acquire exclusive ownership before accessing management data for the buffer storage, the examiner did not explain how such a MUTEX mechanism controls the producer (data source) to control the amount of data stored in a queue.

Applicant's arguments are not persuasive. The MUTEX, of course, does not provide mutually exclusive access to a particular buffer in the buffer storage 372 (queue) itself but to the management data structure that controls the particular buffer. Thus, the scheduler's MUTEX controls access to management data structures such as buffer masks, which controls the producer (data source) to control the amount of data stored in the queue (as discussed in the prior rejection and below). As the examiner previously explained, the MUTEX 374 controls the producer (data source) by placing requirements on the producer to acquire ownership before accessing and modifying management data for buffer storage. Forcing the producer (data source) to meet an ownership requirement before modifying management data structure that affects how much data

is written to the buffer (queue) by the producer would seem to be exercising a restraining or directing influence on (i.e., controlling) on the producer before the producer controls the amount of data the producer writes to the queue. *For example, if the scheduler (storage manager) does not let the producer (data source) acquire the ownership and the ability to modify the management data structure that controls buffer access, then the producer will write no amount of data to the buffer (queue) in accordance with the management structure.* Thus, the MUTEX allows the scheduler (storage manager 350) to control the producer (data source), which in turn controls the amount of data written to the queue (buffer). Again, the applicant uses the unclaimed term "essentially" when concluding that the scheduler must essentially control the amount of data stored in the queue via control of the data source. Although the MUTEX essentially controls the amount of data stored in the queue via the data source (as discussed above), the claims do not require this (also as discussed above).

On page 8 of the brief, the applicant argues that the mailslots do not provide a means for the scheduler (storage device 350) to control the consumer, which in turns controls the amount of data consumed from the queue. Specifically, the applicant argues that the examiner's reliance on the mail slots 340₀-340_N of the IPC channel 330 is misplaced because the mail slots are not part of, or controlled by the storage manager 350 but rather controlled by the producer to notify consumers of the arrival of an information stream.

Applicant's arguments are not persuasive. Each mail slot contains an index, which represents the location of the information stream stored in a particular buffer (queue) in the

storage manager 350 (col. 7, lines 11-13). When the consumers receive a notification that an information stream has arrived, they map this location into their address in order to read an amount of data from the buffer (queue) (col. 7, lines 13-18). Thus, the storage manager (scheduler) exercises a restraining or directing influence (i.e., controls) the consumer by the use of the mail slots. Once the consumer gets the address of the buffer to read the data, it then reads an amount of data from the queue (buffer). *That is, if the scheduler (data source) does not provide a location of buffer (queue) containing the information stream via the mail slot to the consumer, then the consumer will read no amount of data from the buffer (queue).* Thus, the scheduler controls the consumer, which in turn controls the amount of data read (consumed) from the queue. Thus, the scheduler essentially controls the amount of data consumed (read) from the queue via control of the consumer although the claims do not require this interpretation (as discussed above).

On page 8 of the brief, the applicant argues that the buffer masks 364 are nothing more than bit masks that indicate which clients or clients have reference to a given buffer and thus fail to control the consumer to control the amount of data consumed from the queue.

Applicant's arguments are not persuasive. Forcing the client (consumer) to comply with a bit mask indication (as provided by the scheduler) that the buffer is available before the client can use the buffer is a concrete example of the scheduler exercising a restraining or directing influence on the consumer. Further, the once the consumer has a reference to a particular buffer (queue) in the buffer storage area, it can then start reading amounts of data from that queue. *That*

is, if the scheduler's bitmask does not provide the proper indication to the consumer that the buffer (queue) is available, then no amount of data will be read by the consumer from that buffer (queue). Thus, buffer mask indirectly allows the scheduler (storage manager 350) to control the amount of data consumed (read) from the queue via control of the consumer although the claims do not require this interpretation (as discussed above).

Claim 3

On page 9 of the brief, the applicant argues that Woodring does not disclose that the "storage manager (350) of the IPC channel (330) controls or otherwise manages a second queue, much less that the producer (310) consumes data from a second queue managed by the storage manager (350)."

Applicant's arguments are not persuasive. The producer 310 referenced by the examiner in the claim 3 rejection is part of Fig. 4., where Fig. 4 "is generic to any producer-consumer relationship, whether it is local or remote (col. 6, lines 8-17) (emphasis added). Thus, producer 310 would be generic to the producer-consumer relationship illustrated in the video-conferencing arrangement of Fig. 1, where you have a local producer (camera 42 and microphone 62) at PC 12, a remote producer at PC 11 (camera 41 and microphone 61), and a local consumer at PC 12 to receive the streaming data from the remote producer of PC 11. Therefore, PC 12 (part of producer 310) would be associated with a first queue for locally produced data and a second queue for consuming remotely produced data from PC 11. Otherwise if producer 310 was insufficient to handle the duties illustrated in Fig. 1, then Fig. 4 would not be generic to any

producer-consumer relationship whether local or remote as required by Woodring. Similarly, the storage manager 350 (scheduler) of Fig. 4 would be generic to Fig. 1, and thus manage any producer-consumer queues (a concept previously discussed in the claim 1 rejection), including the first and second queues discussed above.

Even if one were to ignore the express teaching of Woodring that Fig. 4, which includes the producer 310 and storage manager 350 (scheduler), is generic to any producer-consumer relationship whether local or remote, it should be noted that each local data source (e.g., the PC comprising video camera 41) also comprises a local consumer 242 received from the same IPC channel 230 that the local data source transmits video on (via video producer 224) (Fig. 3). Therefore, Woodring necessarily discloses a first queue associated with the remote consumers 262 and a second queue associated with the local data source's local consumer 242, where both queues are managed by the same storage manager (scheduler) associated with IPC channel 230.

Claims 4-6, 17, 18, 24, and 25

On page 10 of the brief, the applicant argues that the priority request registration cannot be related to the desired sample rate at which the consumers are notified of arrived streaming data. Further, even if the sample rate equates to priority, then the sampling rate specifiers are processed and accessed by the producer and thus not registered by the consumers.

Applicant's arguments are undercut by applicant's specification and are thus not persuasive. For example, applicant admits that one example of "prioritizing and slowing down

different consumers" is to "slow down network traffic towards fast consumers" (page 27, lines 10-20). Slowing down network traffic to the consumers relates directly to Woodring's sampling rate, which controls the speed (rate) in which the consumers receive data from the network. Therefore, examiner's interpretation of priority is consistent with applicant's specification. Further, the desired sampling rate is part of a consumer-side interface 450 specification provided to (i.e., registered with) the IPC channel 330 (Fig. 5) (see also the applicant's arguments on page 10, which describe the registration process).

Claims 8, 20, and 27

On pages 10 and 11 of the brief, the applicant argues that Woodring fails to disclose or suggest a scheduler that prioritizes data consumption for the first queue based on an amount of unread data of each of the plurality of consumers because the scheduler manages consumption of the queue by prioritizing and slowing down the different consumers. Woodring only teaches that the multiple consumers "read data out of the buffer simultaneously without holding any memory locks of the buffer storage while reading data."

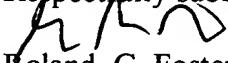
Applicant's arguments are not persuasive. Woodring clearly discloses prioritizing data consumption based on the amount of unread data. Each buffer (queue) holds an amount of unread data represented by the buffer size 368 (col. 7, lines 32-35). The consumers in turn read out this amount of unread data (col. 7, lines 13-35). As discussed in the last Office action, only when the last consumer assigned to a particular buffer releases the buffer after reading this amount of unread data will the buffer mask be set to zero (cleared), thereby allowing the buffer

to become available for other data consumption (e.g., different consumers) (see e.g., col. 7, line 60 – col. 8, line 5). Therefore, the storage manager (scheduler) prioritizes data consumption by keeping the buffer allocated to a consumer based on an amount of unread data (buffer size) being present in the queue (col. 7, lines 13-35). That is, if there is an amount of unread data in the buffer, then the scheduler will prioritize data consumption by keeping the buffer allocated to the consumer. The applicant cannot reasonably argue by plain meaning that holding a buffer in reserve for use by a consumer based on the presence of an amount of unread data in the queue is not prioritizing buffer usage to the advantage of that consumer. The absence of memory locks is not relevant because the buffer mask represents an improved mechanism to prioritize buffer usage without the use of memory locks (col. 8, lines 1-5).

For the above reasons, it is believed that the rejections should be sustained.

Art Unit: 2645

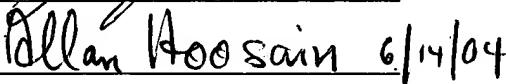
Respectfully submitted,


Roland G. Foster
Primary Examiner
Art Unit 2645

Roland G. Foster (Primary Patent Examiner): 
June 14, 2004

Conferees:

Fan Tsang (Supervisory Patent Examiner): 
Fan Tsang

Allan Hoosain (Primary Patent Examiner): 
Allan Hoosain 6/14/04

Frank Chau Esq
F Chau & Associates LLP
1900 Hempstead Turnpike
Suite 501
East Meadow, NY 11554